

WHAT IS CLAIMED IS:

1. A method for detecting transmission of potentially unwanted e-mail messages, comprising:
 - receiving a plurality of e-mail messages;
 - generating hash values, as generated hash values, based on one or more portions of the plurality of e-mail messages;
 - determining whether the generated hash values match hash values associated with prior e-mail messages; and
 - determining that one of the plurality of e-mail messages is a potentially unwanted e-mail message when one or more of the generated hash values associated with the one of the plurality of e-mail messages match one or more of the hash values associated with the prior e-mail messages.
2. The method of claim 1, wherein the generating hash values includes:
 - performing a plurality of hashes on a plurality of variable-sized blocks of a main text of the plurality of e-mail messages.
3. The method of claim 1, wherein the generating hash values includes:
 - performing a plurality of hashes on a plurality of fixed-sized blocks of a main text of the plurality of e-mail messages.

4. The method of claim 1, wherein the generating hash values includes:
performing a plurality of hashes on a main text of the plurality of e-mail messages using a plurality of different hash functions.

5. The method of claim 1, wherein the generating hash values includes:
performing a plurality of hashes on a main text of the plurality of e-mail messages using a same hash function.

6. The method of claim 1, wherein the generating hash values includes:
attempting to expand an attachment of the plurality of e-mail messages, and
hashing the attachment after attempting to expand the attachment.

7. The method of claim 1, wherein the generating hash values includes:
performing a plurality of hashes on a plurality of variable-sized blocks of an attachment of the plurality of e-mail messages.

8. The method of claim 1, wherein the generating hash values includes:
performing a plurality of hashes on a plurality of fixed-sized blocks of an attachment of the plurality of e-mail messages.

9. The method of claim 1, wherein the generating hash values includes:
performing a plurality of hashes on an attachment of the plurality of e-mail messages
using a plurality of different hash functions.

10. The method of claim 1, wherein the generating hash values includes:
performing a plurality of hashes on an attachment of the plurality of e-mail messages
using a same hash function.

11. The method of claim 1, further comprising:
comparing the generated hash values to hash values corresponding to known unwanted e-mails.

12. The method of claim 11, wherein the known unwanted e-mails include at least one
of e-mails containing a virus, e-mails containing a worm, and unsolicited commercial e-mails.

13. The method of claim 1, wherein the generating hash values includes:
hashing at least one of a main text and an attachment to generate one or more first hash
values, and
hashing a concatenation of first and second header fields to generate a second hash value.

14. The method of claim 13, wherein the first and second header fields include a From header field and a To header field.

15. The method of claim 13, wherein the determining whether the generated hash values match hash values associated with prior e-mail messages includes:

determining a first suspicion count based on a number of the hash values associated with the prior e-mail messages that match the one or more first hash values, and

determining a second suspicion count based on a number of the hash values associated with the prior e-mail messages that match the second hash value.

16. The method of claim 15, wherein the determining that one of the plurality of e-mail messages is a potentially unwanted e-mail message includes:

determining that the one of the plurality of e-mail messages is a potentially unwanted e-mail message when the first suspicion count is significantly higher than the second suspicion count.

17. The method of claim 1, further comprising:

taking remedial action when the one of the plurality of e-mail messages is a potentially unwanted e-mail message, the taking remedial action including at least one of:

discarding the one of the plurality of e-mail messages,
bouncing the one of the plurality of e-mail messages,
marking the one of the plurality of e-mail messages with a warning,
subjecting the one of the plurality of e-mail messages to a virus or worm detection
process,
creating a notification message, and
generating a suspicion score for the one of the plurality of e-mail messages and using the
suspicion score to identify further processing for the one of the plurality of e-mail messages.

18. The method of claim 1, further comprising:

generating a suspicion score for the plurality of e-mail messages based on a result of the
determination of whether the generated hash values match hash values associated with prior e-
mail messages; and

taking remedial action when the one of the plurality of e-mail messages is a potentially
unwanted e-mail message, the taking remedial action including:

determining whether a newly received e-mail message exceeds a mail quota,
identifying an earlier-received e-mail message with a highest suspicion score,
determining whether the suspicion score of the newly received e-mail message is lower
than the suspicion score of the earlier-received e-mail message when the newly received e-mail
message exceeds the mail quota,

deleting the earlier-received e-mail message when the suspicion score of the newly received e-mail message is lower than the suspicion score of the earlier-received e-mail message, and
storing the newly received e-mail message.

19. The method of claim 1, wherein the generating hash values and the determining whether the hash values match hash values associated with prior e-mail messages are performed incrementally as the plurality of e-mail messages are being received.

20. The method of claim 19, further comprising:
generating a suspicion score for the plurality of e-mail messages based on a result of the determination of whether the generated hash values match hash values associated with prior e-mail messages; and
taking remedial action when the suspicion score of an e-mail message of the plurality of e-mail messages is above a threshold, the taking remedial action including rejecting the e-mail message.

21. The method of claim 20, wherein the rejecting occurs before the e-mail message is completely received.

22. The method of claim 1, further comprising:
comparing the generated hash values to known legitimate mailing lists; and
passing the plurality of e-mail messages without further examination when the generated hash values match one or more of the known legitimate mailing lists.

23. The method of claim 22, wherein the comparing the generated hash values includes:
determining whether the plurality of e-mail messages originated from the known legitimate mailing lists.

24. The method of claim 1, wherein the generating hash values includes:
hashing a main text to generate a first hash value, and
hashing sender-related header fields to generate one or more second hash values.

25. The method of claim 24, wherein the sender-related header fields include at least one of a From header field, a Sender header field, and a Reply-To header field.

26. The method of claim 24, wherein the determining whether the generated hash values match hash values associated with prior e-mail messages includes:

determining a first suspicion count based on a number of the hash values associated with the prior e-mail messages that match the first hash value, and

determining one or more second suspicion counts based on a number of the hash values associated with the prior e-mail messages that match the one or more second hash values.

27. The method of claim 26, wherein the determining that one of the plurality of e-mail messages is a potentially unwanted e-mail message includes:

determining that the one of the plurality of e-mail messages is a potentially unwanted e-mail message when the first suspicion count is higher than the one or more second suspicion counts.

28. The method of claim 1, wherein the generating hash values includes:

hashing a main text of the plurality of e-mail messages to generate a main text hash, and

hashing at least one header field of the plurality of e-mail messages to generate at least one header hash.

29. The method of claim 28, wherein the determining whether the generated hash values match hash values associated with prior e-mail messages includes:

determining whether the main text hash matches a substantially higher number of the hash values associated with the prior e-mail messages than the at least one header hash; and

wherein the determining that one of the plurality of e-mail messages is a potentially unwanted e-mail message includes:

determining that the one of the plurality of e-mail messages is a potentially unwanted e-mail message when the main text hash matches a substantially higher number of the hash values associated with the prior e-mail messages than the at least one header hash.

30. A system for detecting transmission of potentially unwanted e-mails, comprising:
means for observing a plurality of e-mails;
means for hashing one or more portions of the plurality of e-mails to generate hash values, as generated hash values;
means for determining whether the generated hash values match hash values associated with prior e-mails; and
means for determining that the plurality of e-mails are potentially unwanted e-mails when one or more of the generated hash values match one or more of the hash values associated with the prior e-mails.

31. A mail server, comprising:
one or more hash memories configured to store count values associated with a plurality of hash values; and
a hash processor configured to:

receive an e-mail message,
hash one or more portions of the e-mail message to generate hash values, as
generated hash values,
increment the count values corresponding to the generated hash values, as
incremented count values, and
determine whether the e-mail message is a potentially unwanted e-mail message
based on the incremented count values.

32. The server of claim 31, wherein when hashing one or more portions of the e-mail message, the hash processor is configured to perform a plurality of hashes on a plurality of variable-sized blocks of a main text of the e-mail message.

33. The server of claim 31, wherein when hashing one or more portions of the e-mail message, the hash processor is configured to perform a plurality of hashes on a plurality of fixed-sized blocks of a main text of the e-mail message.

34. The server of claim 31, wherein when hashing one or more portions of the e-mail message, the hash processor is configured to perform a plurality of hashes on a main text of the e-mail message using a plurality of different hash functions.

35. The server of claim 31, wherein when hashing one or more portions of the e-mail message, the hash processor is configured to:

attempt to expand an attachment of the e-mail message, and
hash the attachment after attempting to expand the attachment.

36. The server of claim 31, wherein when hashing one or more portions of the e-mail message, the hash processor is configured to perform a plurality of hashes on a plurality of variable-sized blocks of an attachment of the e-mail message.

37. The server of claim 31, wherein when hashing one or more portions of the e-mail message, the hash processor is configured to perform a plurality of hashes on a plurality of fixed-sized blocks of an attachment of the e-mail message.

38. The server of claim 31, wherein when hashing one or more portions of the e-mail message, the hash processor is configured to perform a plurality of hashes on an attachment of the e-mail message using a plurality of different hash functions.

39. The server of claim 31, wherein the hash processor is further configured to compare the generated hash values to hash values corresponding to known unwanted e-mails.

40. The server of claim 39, wherein the known unwanted e-mails include at least one of e-mails containing a virus, e-mails containing a worm, and unsolicited commercial e-mails.

41. The server of claim 31, wherein when hashing one or more portions of the e-mail message, the hash processor is configured to:

hash at least one of a main text and an attachment of the e-mail message to generate one or more first hash values, and

hash a concatenation of first and second header fields of the e-mail message to generate a second hash value.

42. The server of claim 41, wherein the first and second header fields include a From header field and a To header field.

43. The server of claim 41, wherein when determining whether the e-mail message is a potentially unwanted e-mail message, the hash processor is configured to identify the e-mail message as a potentially unwanted e-mail message when the count value corresponding to one or more first hash values is significantly higher than the count value corresponding to the second hash value.

44. The server of claim 31, wherein the hash processor is further configured to take remedial action when the e-mail message is a potentially unwanted e-mail message, when taking remedial action, the hash processor is configured to at least one of:

discard the e-mail message,

bounce the e-mail message,

mark the e-mail message with a warning,

subject the e-mail message to a virus or worm detection process,

create a notification message, and

generate a suspicion score for the e-mail message and use the suspicion score to identify further processing for the e-mail message.

45. The server of claim 31, wherein the hash processor is further configured to:

generate a suspicion score for the e-mail message based on the incremented count values,

determine whether a newly received e-mail message exceeds a mail quota,

identify an earlier-received e-mail message with a highest suspicion score,

determine whether a suspicion score of the newly received e-mail message is lower than the suspicion score of the earlier-received e-mail message when the newly received e-mail message exceeds the mail quota,

delete the earlier-received e-mail message when the suspicion score of the newly received e-mail message is lower than the suspicion score of the earlier-received e-mail message, and

store the newly received e-mail message.

46. The server of claim 31, wherein the hash processor is configured to hash the one or more portions of the e-mail message and increment the count values incrementally as the e-mail message is being received.

47. The server of claim 46, wherein the hash processor is further configured to:
generate a suspicion score for the e-mail message based on the incremented count values,
reject the e-mail message when the suspicion score of the e-mail message is above a
threshold.

48. The server of claim 47, wherein the rejecting occurs before the e-mail message is completely received.

49. The server of claim 31, wherein the hash processor is further configured to:
compare the generated hash values to known legitimate mailing lists, and
pass the e-mail message without further examination when the generated hash values
match one of the known legitimate mailing lists.

50. The server of claim 49, wherein the hash processor is configured to:

determine whether the e-mail message originated from one of the known legitimate mailing lists.

51. The server of claim 31, wherein the hash processor is configured to:
hash a main text of the e-mail message to generate a first hash value, and
hash sender-related header fields of the e-mail message to generate one or more second hash values.

52. The server of claim 51, wherein the sender-related header fields include at least one of a From header field, a Sender header field, and a Reply-To header field.

53. The server of claim 51, wherein when determining whether the e-mail message is a potentially unwanted e-mail message, the hash processor is configured to identify the e-mail message as a potentially unwanted e-mail message when the count value corresponding to the first hash value is higher than the count values corresponding to the one or more second hash values.

54. The server of claim 31, wherein when hashing one or more portions of the e-mail message, the hash processor is configured to:

perform a plurality of hashes on a main text of the e-mail message to generate main text hashes, and

hash at least one header field of the e-mail message to generate at least one header hash.

55. The server of claim 54, when determining whether the e-mail message is a potentially unwanted e-mail message, the hash processor is configured to:

generate a score for the main text based on count values corresponding to the main text hashes and a score for the at least one header field based on the count value corresponding to the at least one header hash, and

identify the e-mail message as a potentially unwanted e-mail message when the score for the main text is substantially higher than the score for the at least one header hash.

56. A method for detecting transmission of unwanted e-mail messages, comprising:
receiving a plurality of e-mail messages; and
detecting unwanted e-mail messages from the plurality of e-mail messages based on hashes of previously received e-mail messages, where multiple hashes are performed on each of the plurality of e-mail messages.

57. A method for detecting transmission of potentially unwanted e-mail messages, comprising:

receiving an e-mail message;

generating a plurality of hash values, as generated hash values, over blocks of the received e-mail message, the blocks including at least two of a main text portion, an attachment portion, and a header portion of the received e-mail message;

determining whether the generated hash values match hash values associated with prior e-mail messages; and

determining that the received e-mail message is a potentially unwanted e-mail message when one or more of the generated hash values associated with the received e-mail message match one or more of the hash values associated with the prior e-mail messages.

58. The method of claim 57, wherein the blocks are variable-sized blocks of the received e-mail message.

59. In a network of cooperating mail servers, one of the mail servers comprising:
one or more hash memories configured to store information relating to hash values corresponding to previously-observed e-mails; and

a hash processor configured to:

receive at least some of the hash values from another one or more of the cooperating mail servers,

store information relating to the at least some of the hash values in at least one of the one or more hash memories,
receive an e-mail message,
hash one or more portions of the received e-mail message to generate hash values, as generated hash values,
determine whether the generated hash values match the hash values corresponding to previously-observed e-mails, and
identify the received e-mail message as a potentially unwanted e-mail message when one or more of the generated hash values associated with the received e-mail message match one or more of the hash values corresponding to previously-observed e-mails.

60. A mail server, comprising:
one or more hash memories configured to store count values associated with a plurality of hash values; and
a hash processor configured to:
receive e-mail messages,
hash one or more portions of the received e-mail messages to generate hash values, as generated hash values,

increment the count values corresponding to the generated hash values, as incremented count values, and generate suspicion scores for the received e-mail messages based on the incremented count values.

61. The server of claim 60, wherein the hash processor is further configured to: maintain a counter corresponding to each of the one or more hash memories, and decrement ones of the count values based on the counter.

62. The server of claim 61, wherein the hash processor is configured to: determine when a value of the counter reaches a threshold, and decrement one of the count values each time another one of the count values is incremented after the value of the counter reaches the threshold.

63. The server of claim 62, wherein the hash processor is further configured to: identify a count value to decrement, determine whether the identified count value is non-zero, and decrement the identified count value when the identified count value is non-zero.

64. The server of claim 63, wherein the hash processor is further configured to:

examine next sequential ones of the count values until a non-zero count value is found when the identified count value is zero, and
decrement the non-zero count value.

65. A method for preventing transmission of unwanted e-mail messages, comprising:
receiving an e-mail message;
generating a plurality of hash values, as generated hash values, over portions of the e-mail message as the e-mail message is being received;
incrementally determining whether the generated hash values match hash values associated with prior e-mail messages;
generating a suspicion score for the e-mail message based on the incremental determining; and
rejecting the e-mail message when the suspicion score of the e-mail message is above a threshold.

66. The method of claim 65, wherein the rejecting occurs before the e-mail message is completely received.